

CLAIMS

What is claimed is:

1. A system comprising:
a first node that provides a source broadcast request for data, the first node being operable to respond in a first manner to other source broadcast requests for the data while the source broadcast request for the data is pending at the first node;
the first node being operable to respond in a second manner to the other source broadcast requests for the data in response to receiving an ownership data response at the first node.
2. The system of claim 1, wherein the ownership data response comprises an indication to the first node that the data associated with the ownership data response comprises migratory data.
3. The system of claim 2, wherein the migratory data comprises a cache ordering point for serializing source broadcast requests for the data, the cache ordering point migrating to the first node from a node that provides the ownership data response.
4. The system of claim 3, wherein the first node is operative to provide an ownership data response to a second node requesting the data, such that the cache ordering point migrates from the first node to the second node.
5. The system of claim 1, wherein the source broadcast request from the first node comprises a source broadcast read request, the first node, when responding in the first manner, provides a first response to the other source broadcast requests for the data indicating that the first node has a conflicting read request for the data.
6. The system of claim 5, further comprising a second node that provides one of the other source broadcast requests for the data and receives the first response from the first node, the second node being operative to fill a shared copy of data received from a third node in response to the one of the other source broadcast requests for the data.

7. The system of claim 5, further comprising a second node that provides one of the other source broadcast requests for the data and receives the first response from the first node, the second node being operative to fill a copy of data received from a home node for the data.

8. The system of claim 1, wherein the first node, when responding in the second manner, provides a second response to the other source broadcast requests for the data indicating that the source broadcast request from the first node is a conflicting request for the data and that migration of the data to the first node is in progress.

9. The system of claim 8, further comprising a second node that provides one of the other source broadcast requests for the data and receives the second response from the first node, the second node being operative to employ a copy of the data received from a third node for only a single use.

10. The system of claim 8, further comprising a second node that provides one of the other source broadcast requests for the data and receives the second response from the first node, the second node being operative to employ a forward progress technique to obtain the data.

11. The system of claim 10, wherein the forward progress technique comprises a forward progress cache coherency protocol.

12. The system of claim 1, wherein the first node employs an invalidate line command to other nodes of the system to remove incorrect copies of the data and any stale copies of the data cached at the other nodes of the system.

13. The system of claim 1, wherein the source broadcast request provided by the first node is broadcast using a source broadcast cache coherency protocol.

14. The system of claim 1, wherein the first node defines a processor having an associated cache, the associated cache of the processor comprising a plurality of cache lines, each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the

respective cache line, the processor being capable of communicating with other nodes of the system through an interconnect, the system further comprising a cache controller associated with the processor, the cache controller being operative to manage data requests and responses for the associated cache of the processor, the cache controller effecting state transitions associated with the data in the associated cache of the processor based on the data requests and responses for the associated cache of the processor.

15. The system of claim 1, wherein the system implements a hybrid cache coherency protocol wherein the first node employs a source broadcast-based protocol to issue the source broadcast request for the data, the first node employing an associated forward progress protocol to reissue a request for the data in response to the request failing in the source broadcast protocol.

16. A multi-processor network comprising:
a source processor node that provides a source broadcast read request for data;
the source processor node issuing an invalidate line command to other processor nodes of the system in response to receiving a data response that transfers a cache ordering point for the data to the source processor node.

17. The multi-processor network of claim 16, wherein the invalidate line command issued by the source processor node removes incorrect cached copies of the data at the other processor nodes of the system and stale copies of the data filled at the other processor nodes of the system.

18. The multi-processor network of claim 16, wherein the source processor node is operative to provide a first conflict response to source broadcast requests for the data from the other processor nodes prior to receiving the data response that transfers the cache ordering point for the data to the source processor node,

the source processor node being operative to provide a second conflict response to at least one source broadcast request for the data from at least one of the other processor nodes in response to the source processor node receiving a conflict response and receiving the data response that transfers the cache ordering point for the data to the source processor node.

19. The system of claim 18, wherein the source processor node provides the first response to the source broadcast requests for the data from the other processor nodes when the source processor node has a pending conflicting read request for the data.

20. The system of claim 19, wherein the other processor nodes receiving the first response from the source processor node are operative to fill a copy of the data received from at least one of the other processor nodes and from system memory.

21. The system of claim 18, wherein the source processor node provides the second response to the source processor node receiving a request for the data that conflicts with the source broadcast request for the data after migration of the data to the source processor node has begun.

22. The system of claim 21, wherein one of the other processor nodes comprises a second processor node that provides a respective one of the other source broadcast requests for the data and receives the second response from the first node, the second processor node being operative to employ a copy of the data received from a third node for a single use.

23. The system of claim 21, wherein one of the other processor nodes comprises a second processor node that provides one of the other source broadcast requests for the data according to broadcast-based protocol, the second processor node being operative to employ a forward progress technique to obtain the data in response to the second response from the first node.

24. The system of claim 23, wherein the forward progress technique comprises a forward progress cache coherency protocol.

25. A system comprising:
means for broadcasting a source broadcast request for data from a first node; and
means for issuing from the first node an invalidate line command to other nodes of the system in response to receiving a conflict response from at least one other node in the system and a data response transferring a cache ordering point for the data to the first node.

26. The system of claim 25, further comprising:

means for providing a first conflict response to source broadcast requests for the data from other nodes prior to receiving a data response transferring a cache ordering point for the data to the first node; and

means for providing a second conflict response to source broadcast requests for the data from other nodes in response to receiving a data response transferring a cache ordering point for the data to the first node.

27. The system of claim 26, further comprising means for filling a shared copy of the data at one of the other nodes in response receiving the first conflict response from the first node.

28. The system of claim 26, further comprising means for filling a copy of the data received from system memory at one of the other nodes in response to receiving the first conflict response from the first node.

29. The system of claim 26, further comprising means for employing a shared copy of the data for a single use at one of the other nodes in response to the one of the other nodes receiving the second conflict response from the first node.

30. The system of claim 26, further comprising means for employing a forward progress technique at one of the other nodes to obtain the data in response to receiving the second conflict response from the first node.

31. A system comprising:

means for broadcasting a source broadcast request for data from a first node;

means for providing a first conflict response to other source broadcast requests for the data from other nodes while the source broadcast for the data is pending at the first node; and

means for providing a second conflict response to the other source broadcast requests for the data from the other nodes after receiving an ownership data response at the first node while the source broadcast for the data is pending at the first node.

32. The system of claim 31, further comprising means for cleaning-up incorrect copies of the data and stale copies of the data filled at other nodes of the system in response to receiving the ownership data response at the first node.

33. The system of claim 31, further comprising means for issuing an invalidate line command to the other nodes of the system in response to receiving the ownership data response at the first node.

34. A method comprising:

migrating a cache ordering point for a line of data from a first node of a system to a second node of a system; and

issuing an invalidate line command for the line of data from the second node to other nodes of the system in response to receiving a conflict response from at least one other node in the system and to the cache ordering point migrating from the first node to the second node.

35. The method of claim 34, further comprising:

providing a first conflict response from the second node to requests for the line of data from the other nodes of the system prior to the cache ordering point migrating from the first node to the second node; and

providing a second conflict response from the second node to requests for the line of data from the other nodes after the cache ordering point migrates from the first node to the second node.

36. The method of claim 35, further comprising:

enabling a shared copy of the line of data to be filled at one of the other nodes of the system in response to receiving the first conflict response from the second node and a data response from at least another node of the system; and

enabling a copy of the line of data received from system memory to be filled at one of the other nodes of the system in response to receiving the first conflict response from the second node.

37. The method of claim 35, further comprising enabling a shared copy of the line of data to be filled at least one of the other nodes of the system for a single use by the at least

one of the other nodes of the system in response to receiving the second conflict response from the first node.

38. The method of claim 35, further comprising employing a forward progress technique at the other nodes to fill the cache line in response to receiving the second conflict response from the first node.

39. A method comprising:

providing a first conflict response from a first node to source broadcast requests for data from other nodes while a source broadcast request for the data is pending at the first node; and

providing a second conflict response from the first node to the other source broadcast requests for the data from the other nodes in response to receiving a conflict response and an ownership data response at the first node.

40. A computer system comprising a plurality of nodes, the plurality of nodes employing a cache coherency protocol operative to migrate a cache ordering point for a line of data from a target node to a source node in response to a source broadcast read request for the line of data issued by the source node, the source node being operative to invalidate the line of data at other nodes of the computer system in response to receiving a conflict response and migratory data to the source broadcast read request.